

Using sentence connectors for evaluating MT output

Eric M. Visser⁰ and Masaru Fuji

Fujitsu Laboratories Ltd., Media Integration Laboratory
4-1-1 Kamikodanaka, Nakahara-ku, Kawasaki 211, Japan
{eric|fuji}@flab.fujitsu.co.jp

Abstract

This paper elaborates on the design of a machine translation evaluation method that aims to determine to what degree the meaning of an original text is preserved in translation, without looking into the grammatical correctness of its constituent sentences. The basic idea is to have a human evaluator take the sentences of the translated text and, for each of these sentences, determine the semantic relationship that exists between it and the sentence immediately preceding it. In order to minimise evaluator dependence, relations between sentences are expressed in terms of the **conjuncts** that can connect them, rather than through explicit categories. For an n -sentence text this results in a list of $n - 1$ sentence-to-sentence relationships, which we call the text's **connectivity profile**. This can then be compared to the connectivity profile of the original text, and the degree of correspondence between the two would be a measure for the quality of the translation.

A set of "essential" conjuncts was extracted for English and Japanese, and a computer interface was designed to support the task of inserting the most fitting conjuncts between sentence pairs. With these in place, several sets of experiments were performed.

1 Background

Evaluation of MT results is generally tackled on a very detailed, linguistic-technical level. Typically, a test set of sentences is prepared each of which is carefully designed to ascertain whether the MT system can handle a certain grammatical phenomenon — e.g. (Isahara, 1995). Other

methods may concentrate on word choice, consistency in terminology, PP attachment, dependency relations, or other specific grammatical or lexical aspects. While such evaluation methods are certainly necessary and useful for the MT developer, they do not necessarily give us a reliable indication of user satisfaction.

Especially now that MT systems are becoming widely available on the home user market and coming within the casual user's reach, MT developers need to pay more attention to this aspect. Casual users might just not care all that much about grammatical correctness: as long as they can *understand* the output, they might be satisfied with the system. Moreover, such users are not likely to judge the system on a sentence-by-sentence basis: rather, they will be interested in the understandability of the text as a whole. The flurry of integrated WWW-browsers cum MT systems¹ to hit the (Japanese) market recently has added to the plausibility of this scenario.

We conclude then that an MT evaluation method is called for which concentrates on whole texts rather than on single sentences, and which judges meaning and readability rather than grammar. In addition, we specify that evaluation results should be reproducible and evaluator-independent (to a reasonable degree at least), and quantifiable. These additional requirements are necessary to ensure that results obtained at different times and/or by different evaluators (preferably using different texts) are comparable.

In (Su et al., 1992) an interesting alternative evaluation method is proposed, in which the discrepancy is measured between raw MT output and the post-edited result. This method does work on whole texts, and could conceivably be adapted to judge meaning and readability (by adequately instructing the post-editors); then again in "browsing" applications post-editing is not the norm, and it may be difficult to attain a good approximation of "browsable" MT output. In this paper, we try a different approach.

⁰The first author is currently at ATR Interpreting Telecommunications Research Laboratories; current e-mail address is (eric@itl.atr.co.jp).

¹These allow you to read English WWW-pages in Japanese, preserving the original page layout.

2 Outline of the evaluation method

2.1 Compare salient properties

To test whether the meaning of a translated text has come across, one could simply ask the evaluators questions about the translated text, or have them summarise it. Such methods however are either costly (for each new text a new set of questions will have to be devised) or hard to quantify objectively, or even both.

The method we will adopt involves constructing a profile of both the original and the translated text in terms of some salient semantic or pragmatic property of its constituent sentences. These profiles can then be compared to give an indication of translation quality: if we assume that the original text's profile is "perfect", then the degree to which the profile of the translated text resembles the perfect profile will correspond (in theory at least) to the quality of the translation. This approach assumes that the number and order of sentences are invariant in translation; luckily, for MT systems, this is almost always true.

As for the salient property to be used in the profile, we settled on meaning relations of single sentences with previous text: this property seemed to us to be both fairly discriminating and implementable. In summary, a profile will be an ordered list of meaning relations x_i , $i = 2 \dots n$ which describe the relation of sentence i with what came before. Moreover, the target of each relation is taken to be the previous sentence, i.e. sentence $i-1$ (see § 4 for further discussion).

2.2 Avoid contrived definitions

A set of sentence-to-sentence relation categories will then have to be designed and defined; but the wide variety of proposed methods and solutions (see (Hovy and Maier, 1993) for an overview) suggests that this is not an easy task. Indeed, the problem with categories and definitions is that the evaluator will always have to depend to a certain extent on his own personal understanding of these definitions; and the more categories there are, the greater the chance that their definitions will not always be clear and fixed in his mind. This naturally has a deleterious effect on the reliability and universality of evaluation results.

We will get back to the design problem later, but with respect to the definition problem, our solution was to simply hide the definitions. We have sought to accomplish this by instructing the evaluator to link sentences linguistically; more specifically, we have opted to instruct the evaluator to choose a **conjunct**² to be inserted between every pair of consecutive sentences. The conjuncts

themselves may be divided into categories, but these can remain hidden from the evaluator. This approach hinges on the hope that straight linguistic knowledge comes more naturally to people and is less susceptible to person-to-person differences than contrived meaning categories.

2.3 Standardise thinking methods

Small-scale preliminary experiments (on paper) showed that in spite of the above refinements, evaluator differences were still larger than seemed reasonable. We surmised that this was due to differences in work methods (or thinking methods), and that therefore these needed to be equalised a little more. We decided on two countermeasures.

Recognising that the class of conjuncts was too large for the evaluator to encompass at a glance, we decided to implement an interactive Q&A interface on the computer in order to gradually guide the evaluator to the optimal choice of a conjunct. Obviously this opens a whole new can of worms, in that the interface has to be designed (the kind and order of questions etc.); we will get back to that later (in § 4).

The other step was to instruct the evaluator to extract the **topic** and **comment** of the sentence under consideration. Both topic and comment were only loosely defined: in truth the topic and comment are not important as such, rather their extraction was intended as a means to force the evaluator to get a clearer picture of the meaning of the sentence under consideration (though we did not tell them this).

3 Basic assumptions

At this point, it is useful to look back at the design considerations outlined above and to clarify exactly what assumptions on sentences and relations underlie them. With a little luck, our results can provide some support for these assumptions.

The first of our assumptions is that it is always possible to make explicit the relationship of a sentence to what has come before using a conjunct. The conjunct may be present in the sentence, but even if it is not, it can be added in a linguistically satisfactory way. We also assume that the assignment of acceptable conjuncts is reader-independent to a large degree.

We assume that conjuncts (which form a closed class) can be divided into a limited number of categories that are meaningful in terms of expressing the semantic relationship between sentences.

Yet another assumption is that the meaning relationships between sentences of a text combine to form a characteristic feature (a profile) of that text, and that this profile needs to be preserved in translation. Moreover, the ease with which this profile can be discerned in the translated text is assumed to be related to the readability or understandability of the text as a whole.

²A subclass of the adverbs, cf. (Quirk et al., 1985) pg. 631-. For languages that do not recognise this class, surrogates can be concocted: for Japanese, a mixture of conjunctions and conjoining adverbs.

4 The implementation

A prototype was implemented on a Macintosh computer using HyperCard. The evaluation process is made up of the following steps, which have to be executed for every sentence in the text.

1. Extract the topic(s) and comment(s) of the sentence under consideration.
2. If there is more than one topic/comment pair, order the pairs as seems best and determine (using the same method as for sentences) which conjuncts fit best between the pairs.
3. Determine through a dialog with the system which conjunct fits best at the start of the sentence under consideration.

A backtrack function was implemented which allowed the subjects to come back on decisions made earlier in the dialog. The prototype keeps a very detailed log of what the evaluator does exactly. Without going into technical details, the following were the main tasks in the implementation.

Categorising the conjuncts

Our first categorisation of conjuncts was based on information concerning conjuncts and rhetorical structures that we patched together from authoritative grammars for English (Quirk et al., 1985), Japanese (Martin, 1975) e.a. We came up with 9 categories; in a later redesign we took the conjuncts themselves as our starting point and, by tracing crossreferences in dictionaries, were able to reduce the initial number of ± 220 to 32 “basic” conjuncts, divided over 11 categories.

Assisting topic/comment extraction

Frankly we have been unable to find a foolproof method, and have settled for user-requested online help cued on linguistic aspects of the sentence.

Defining the scope of meaning relations

We have established above that meaning relations hold between consecutive sentences; this is however not self-evident. A sentence may relate to a more remote sentence (i-5, for instance), or to a block of sentences; see (Kurohashi and Nagao, 1995) for a more plausible model. We found however that an online computer interface that would allow the user to specify the target of a relation to this extent would become prohibitively complicated. The evaluator’s task would involve so much juggling with relations and attaining such a deep understanding of the text that it would in the end have a negative effect on the reproducibility and evaluator-independence of the results.

Designing the dialog

We believe that this is a trial-and-error process which will have to be guided by the outcome of experiments; more about this will follow below.

	A	B	C	D
mean	0.43	0.46	0.37	0.32
time (m:s)	13:27	14:16	12:32	41:23
backtracks	11.8	9.8	4.4	2.6

Table 1: Results of the first experiment

5 The experiments

We decided that experiments needed to establish three qualities of this system.

Evaluator-independence Given a text in one language, different evaluators should produce the same connectivity profile.

Language-independence Given a “perfectly” translated text, its connectivity profile should turn out the same as that of the original.

Quantifiability Given translations of varying quality, the degree of correspondence in the connectivity profiles must be shown to correspond to the quality of the translation.

But first we conducted a preliminary experiment.

5.1 Experiments with the dialog

Our first experiments (Japanese only) concerned the conjunct-determining dialogs. We implemented 3 interfaces, each comprising the same 61 conjuncts spread over 9 categories: one (A) based on categories (the subjects got a list of categories in the first screen, and if they clicked one they got the conjuncts in that category on the second screen); one (B) based on the conjuncts themselves (the subjects just got the whole list of conjuncts, spread over a couple of screens, without elaboration); and one (C) with questions (3 answers to choose from on the first screen, one of these leads to a second question with 4 answers, all other links lead to sets of conjuncts).

Subjects were assigned an interface, given a 9-sentence text and asked to connect the sentences, without however performing topic/comment extraction. A fourth group was asked to use interface C, but also to extract topic and comment before connecting the sentences (D). The results are given in table 1. The mean of the evaluators’ choices was computed by transforming the results into numbers (if 7 out of 10 evaluators chose category X, 2 chose Y, and 1 Z, then this would result in the values {1 1 1 1 1 1 2 2 3}), and inputting these numbers into the following formula.

$$\mu_2^e = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2$$

We might add that subjects using interfaces A and B were more likely to choose “safe” (ambiguous, vague) conjuncts such as ‘soshite’ (and then), and also — for what it’s worth — complained more.

	A (14)	B (13)	C (7)	D (7)
mean(cat)	0.52	0.60	0.89	0.69
mean(con)	1.99	2.66	2.20	1.76
time (m:s)	15:44	19:38	13:40	14:42
backtracks	4.6	9.8	6.7	6.1

NB: (C+D) mean(cat) = 1.65, mean(con) = 4.91.

Table 2: Experiment results for the various texts

	A+B	A+C	A+D	A+C+D
mean(cat)	0.73	0.98	1.02	1.59
mean(con)	4.62	4.50	4.29	7.91

Table 3: Combined experiment results

To be quite honest this experiment was too small in scale to allow scientific conclusions (20 people participated), but we went ahead anyway and concluded that a) the project showed promise, b) interface C was the way to go, c) topic/comment extraction was important, but d) it was also costly (took three times as long!) so we'd stick to the 'lazy' evaluation for further experiments.

5.2 Validation experiments

For the second set of experiments, we designed identical interfaces for English and Japanese. There was only one question, with 6 answers, and all of these led to a screen with conjuncts to choose from, never more than 8 on a screen. The set of conjuncts was designed to be minimal (no redundancies, no ambiguous conjuncts); there were 32 of them, spread over 11 categories (cf. § 4).

An original English text was chosen (A); then a "perfect" (but aligned) Japanese translation was produced (B); and finally two "less-than-perfect" translations were contrived (C was raw MT output, D was output from a tuned MT system the understandability of which had been determined by independent experiments to be halfway between B and C — level 3 in (Fuji, 1996)). The sizes of the subject groups are given in table 2 between parentheses. Distribution means were computed both for categories and for conjuncts.

6 Discussion

The category means basically follow expectations. Those of C and D come out a bit low, but the combined mean for C+D suggests that this may be partly due to the size of the sample. The conjunct mean of B is very high; it is not clear why. It must be noted that the evaluators were totally untrained; in the context of the intended use of this method, requiring a certain level of training seems acceptable and this would surely bring results closer to the goal of evaluator independence.

However, we also observed several instances where the choice of a conjunct was dictated by the evaluator's prior knowledge (or lack of it) of the subject area; this is a discrepancy we cannot resolve.

The cross-linguistic category mean for A+B is significantly lower than that of A+C and A+D. The conjunct mean is rather high: this is probably due to the unexplained high conjunct mean for B. The conjunct means of A+C and A+D seem to correlate with the number of unintelligible sentences in the machine-translated texts. Again the means of A+C+D are fairly enormous, indicating that size is still a factor.

A rather unsettling result, however, was that the most-chosen sentence connector was identical across texts for almost each of the sentence pairs. This suggests that reducing evaluator dependence will lower *all* means, which would defeat the purpose of this research.

In conclusion, we feel justified in hoping that the goals of evaluator-independence and language-independence are reachable through judicious tuning of the current system. The project has also been successful in that it has yielded a wealth of interesting data about sentence connections. It is doubtful however that the approach will give a useful indication of translation quality.

References

- Masaru Fuji. 1996. Experiments to evaluate the reading comprehension of E-J machine-translated texts. In *Proceedings of the 2nd Annual Meeting of the Ass. for NLP (Japan)*, pages 21–24. In Japanese.
- Eduard H. Hovy and Elisabeth Maier. 1993. Par-simonious or profligate: How many and which discourse structure relations? Technical report, Information Sciences Institute, University of Southern California.
- Hitoshi Isahara. 1995. JEIDA's test-sets for quality evaluation of MT systems: Technical evaluation from the developer's point of view. In *Proceedings of the MT Summit V*. Luxembourg.
- Sadao Kurohashi and Makoto Nagao. 1995. Automatic detection of discourse structure by checking surface information in sentences. *Journal of the Ass. for NLP*, 1(1):3–20. In Japanese.
- Samuel E. Martin. 1975. *A reference grammar of Japanese*. Yale University Press, New Haven.
- Randolph Quirk, Sidney Greenbaum, Geoffrey Leech, and Jan Svartvik. 1985. *A comprehensive grammar of the English language*. Longman, London.
- Keh-Yih Su, Ming-Wen Wu, and Jing-Shin Chang. 1992. A new quantitative quality measure for machine translation systems. In *Proceedings of COLING-92*, pages 433–439.